



EUROPEAN PATENT APPLICATION

(43) Date of publication:
12.12.2018 Bulletin 2018/50

(51) Int Cl.:
G10L 15/30^(2013.01) G10L 15/32^(2013.01)

(21) Application number: **18177044.7**

(22) Date of filing: **11.06.2018**

(84) Designated Contracting States:
**AL AT BE BG CH CY CZ DE DK EE ES FI FR GB
GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO
PL PT RO RS SE SI SK SM TR**
Designated Extension States:
BA ME
Designated Validation States:
KH MA MD TN

(71) Applicant: **SoundHound, Inc.**
Santa Clara, CA 95054 (US)

(72) Inventor: **Mont-Reynaud, Bernard**
Sunnyvale, CA California 94087 (US)

(74) Representative: **Smith, Jeremy Robert et al**
Kilburn & Strode LLP
Lacon London
84 Theobalds Road
London WC1X 8NL (GB)

(30) Priority: **09.06.2017 US 201715619304**

(54) **DUAL MODE SPEECH RECOGNITION**

(57) A dual mode speech recognition system sends speech to two or more speech recognizers. If a first recognition result is received, whose recognition score exceeds a high threshold, the first result is selected without waiting for another result. If the score is below a low threshold, the first result is ignored. At intermediate values of recognition scores, a timeout duration is dynamically determined as a function of the recognition score.

The timeout duration determines how long the system will wait for another result. Many functions of the recognition score are possible, but timeout durations generally decrease as scores increase. When receiving a second recognition score before the timeout occurs, a comparison based on recognition scores determines whether the first result or the second result is the basis for creating a response.

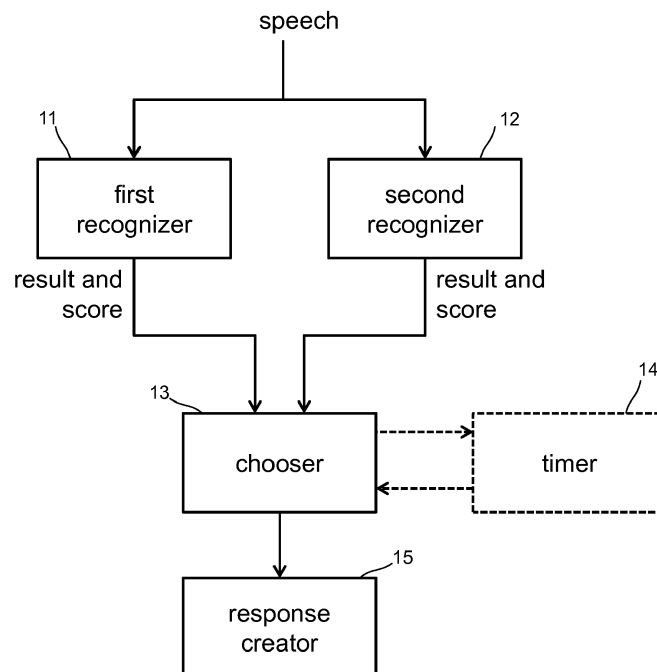


FIG. 1

Description

FIELD OF THE INVENTION

[0001] The present invention is in the field of automated speech recognition (ASR), and, more specifically, use of multiple recognizers.

BACKGROUND

[0002] An objective of some speech recognition systems is to perform recognition as accurately as possible. An objective of some speech recognition systems is to respond with the most useful possible results. An objective of some speech recognition systems is to respond robustly in environments with failure-prone connections. An objective of some speech recognition systems is to distribute the processing workload between computer processors or geographical locations such as server farms.

SUMMARY OF THE INVENTION

[0003] An invention is set out in the independent claims.

[0004] The present disclosure is directed to embodiments of systems, methods, and non-transitory computer readable media that perform dual mode speech recognition. Various embodiments respond to user speech immediately if a result is of sufficient quality. Quality is measured by a recognition score. Various embodiments respond early if the first result is useful, as measured by the quality score, and vary the latency as a function of the quality of the result. Various embodiments use timeout events whose duration varies with quality: a low quality result suggests waiting longer for a higher quality result. Various embodiments ignore early results if they are below an acceptable level of quality, and respond with a later result or an error if no second result is received before a timeout occurs.

[0005] Some embodiments have asymmetrical recognizers, such as one that responds more quickly and one that responds with more accurate or more useful results. For example, some mobile phones perform speech recognition both locally and over a wireless Internet connection. Some earpiece headsets perform speech recognition in the headset, but also in a phone connected over a personal area network.

[0006] Some embodiments are Internet-connected automobiles that respond, if possible from a remote server, which has access to useful dynamic data such as weather and traffic conditions, but responds from a local recognizer when the automobile is in a location that has no wireless network connectivity.

[0007] Some embodiments are power-sensitive systems-on-chip that use low power processors for recognition in a typical mode, but wake up a high performance processor if needed to provide better results.

[0008] Some embodiments use server-based dual mode recognition, and send speech to more than one server with symmetrical recognizers that have different ping latencies or different presence of local data.

[0009] Some embodiments send speech to multiple remote recognizers if accessing different recognizers incur different costs. In such case, it may be advantageous to send spoken utterances to the less expensive one, compare a resulting recognition score to a threshold, and, if the recognition score is below the threshold, send the spoken utterance to a second recognizer.

[0010] Some embodiments perform recognition on delimited spoken queries, such as the speech between a detected wake-up phrase and a detected end-of-utterance. Some embodiments perform recognition continuously, typically at periodic intervals called frames, such as every 10 msec. Some embodiments perform speech recognition incrementally.

[0011] Various embodiments quantify the quality of results using various appropriate techniques. As part of speech recognition, some embodiments compute hypotheses and probability scores for phonemes, phonetic sequences, word sequences (transcriptions), grammatically correct sentences (parses), and meaningful interpretations. Recognition scores, in various embodiments, are based on a probability score alone or a combination of such probability scores.

[0012] The inventor has appreciated that the objectives discussed above can conflict with another objective, which is to respond to users with useful results as quickly as possible. That is, with low latency. The present disclosure achieves both speed and accuracy by immediately accepting the first-received result if it is of sufficient quality, while otherwise waiting for a second result if the first-received result is not of sufficient quality. If the second result is of higher quality then the overall quality is improved, but the delay caused by waiting only takes place if the first-received result was insufficient.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013]

FIG. 1 illustrates components of a dual mode speech recognition system, according to some embodiments.

FIG. 2A illustrates a timeline of events for choosing between two results, according to some embodiments.

FIG. 2B illustrates a timeline of events for choosing a first result because of receiving a timeout before receiving a second result, according to some embodiments.

FIG. 3 illustrates choosing between a first and second result based on a set according to a first recognition score, according to some embodiments.

FIG. 4 illustrates choosing between a first and second result based on a set according to a first recog-

inition score, but choosing the first result without waiting if the first score exceeds a threshold, according to some embodiments.

FIG. 5 illustrates choosing between a first and second result based on a set according to a first recognition score but choosing the higher scoring result if a second result is received before the timeout occurs, according to some embodiments.

FIG. 6 illustrates choosing between a first and second result based on a set according to a first recognition score, but choosing the first result without waiting if the first score exceeds a high threshold or ignoring the first result if it is below a low threshold, as well as updating a first recognizer, according to some embodiments.

FIG. 7A illustrates several ways to set a timeout duration as a function of a first recognition score, according to some embodiments.

FIG. 7B illustrates choosing immediately or waiting for a second result as a function of a first recognition score, according to some embodiments.

FIG. 8 illustrates a flow chart of choosing a result when its score exceeds a threshold, according to some embodiments.

FIG. 9 illustrates a flow chart of choosing a first result if its score exceeds a threshold, and otherwise, sending speech to a second recognizer, according to some embodiments.

FIG. 10 illustrates a flow chart of setting a timeout duration as a function of a recognition score, according to some embodiments.

FIG. 11 illustrates a flow chart of setting a timeout duration as a function of a recognition score and choosing a first result if a second result is not received before the timeout occurs, according to some embodiments.

FIG. 12 illustrates a mobile phone embodiment.

FIG. 13 illustrates a server embodiment.

FIG. 14A-14C illustrates an earpiece headset embodiment.

FIG. 15A-15B illustrates non-transitory computer readable media embodiments.

FIG. 15C-15D illustrates a computer chip embodiment.

FIG. 16 illustrates a block diagram of a system-on-chip embodiment.

FIG. 17 illustrates a block diagram of a server system embodiment.

DETAILED DESCRIPTION

Terminology as used herein

[0014] Spoken utterances are delimited segments of speech, typically comprising multiple words. In various embodiments, they are initiated by a wake-up phrase or a UI action such as clicking or tapping, and terminated by detection of an end-of-utterance event or a UI action

such as tapping or releasing a button.

[0015] Recognizers are hardware- or software-implemented subsystems, which receive speech (i.e. speech audio data) and return recognition results with associated scores. The form and the nature of results varies widely across embodiments, but can include a text transcription, information requested by the speech, or representations of user intents, as data structures represented in JavaScript Object Notation (JSON) or other equivalent internal or exchange data format.

[0016] Various embodiments respond to users so as to give them a feeling that their speech has effected a desired result. Responses comprise results, but may include other information or actions as appropriate for various embodiments. For example, a spoken user request for a coffee causes a speech-enabled coffee maker to respond with spoken words and to produce a cup of coffee. Results are the basis for embodiments to produce responses. In some embodiments, results are text for a machine to output from a text-to-speech module. In some embodiments, results include text with mark-up meta information and instructions encoded for a machine to process.

[0017] Various recognizer embodiments associate recognition scores with results, and return the scores within results, or prior to results, depending on the embodiment. For example, the recognition score may be presented as included within the results or alternatively as separate from the results. Recognizers produce scores in various appropriate ways that practitioners of the embodiment know.

[0018] Local recognizers are ones present within devices with which users interact directly. Remote recognizers are ones that couple with user devices through means such as networks, cables, or wireless signaling.

[0019] The term timeout can refer to a period (i.e., duration) of time, a point in time, an event, or a stored value, as will be apparent to readers skilled in the art. Various embodiments start a timeout timer counting as soon as they send speech to a recognizer, start a timeout timer upon receiving a first score, or start a timeout timer at the time of any other event, as appropriate.

[0020] A function of a recognition score may be used to determine a timeout duration that is appropriate for a given recognition score. Some embodiments use discrete and some use continuous functions. For many embodiments, a non-increasing function is appropriate.

Various embodiments

[0021] FIG. 1 shows components of a dual mode speech recognition system, according to an embodiment. Speech is received by a first recognizer 11 and a second recognizer 12. Each provides a result and a recognition score. A chooser 13 chooses the result of either the first recognizer or the second recognizer based at least on their associated scores. In some embodiments, chooser 13 also chooses the result from the first recog-

nizer or the second recognizer based on when the results from each were received relative to a timeout event. In such an embodiment, a timeout value may be sent to timer 14, causing the timer to generate a timeout event at a point in time represented by the timeout value. For example, if a time duration is sent to timer 14, the timer may generate a timeout event at a time point determined by the current time plus the timeout value. If a time point is passed to the timer, the timer may generate a timeout event when the current time reaches the specified time point. In another embodiment, a timeout signal may be received by receiving an event notification from another element not determined based on time (and not shown in FIG. 1). The chooser 13 provides the chosen result to response creator 15, which uses the result as a basis for a response.

[0022] FIG. 2A shows a timeline of events that can occur in an embodiment. At time 21, a user device sends speech to a plurality of recognizers. At time 22, chooser 13 receives a score and result from a first recognizer and sets a timeout duration. At time 23, the chooser 13 receives a second result from the second recognizer before the timeout. Accordingly, the chooser 13 considers both results before making a choice for producing a response. FIG. 2B shows an alternative timeline of events. At time 24, the chooser 13 receives a timeout before a second result. Accordingly, the chooser 13 considers only the first result in order to produce a response.

[0023] FIG. 3 shows an embodiment for choosing a first or second result. In some embodiment, the timeout duration may be preconfigured and the timer set before requests are sent to any recognizer. In FIG. 3, the incoming score is associated with a result received from a first recognizer. The request to the second recognizer may have been sent at the same time that the request to the first recognizer was sent or the request to the second recognizer may be sent after the score from the first recognizer is received. The timeout duration for waiting for a result and score from the second recognizer is set to a value depending on the score from the first recognizer. A reason for varying the timeout value for waiting for a result from the second recognizer is that the better the score received from the first recognizer, the lower the probability that the second recognizer will return a score that is higher, and thus, there is less motivation to wait a long time for a result and score from the second recognizer. Conversely, the lower the score received from the first recognizer, the greater the probability that the second recognizer will return a higher score, and thus, more patience in waiting for the second results is justified. Thus, the function for determining the timeout value may be inversely related (such as inversely proportional) to the score (when a higher number score indicates greater confidence in the result).

[0024] In the embodiment of FIG. 3 there is higher confidence in a result received from second recognizer than a result received from a first recognizer. The second recognizer may provide a more confident result because it

has access to a larger database of vocabulary or utilize a more powerful recognition engine. Often, the first recognizer is local to a user's mobile device having relatively limited storage space and computation power as compared to a second recognizer located on a remote server.

[0025] In FIG. 3, the chooser receives a score and sets a timeout duration as a function of the score in Step 34. If the timeout occurs in Step 35 before receiving a second result, the embodiment chooses the first result in Step 36 as the basis for creating a response. Upon receiving a second result before the timeout occurs in Step 35, the embodiment presumes that the second result is superior to the first result and chooses it in Step 37. Such a presumption would be justified, for example, when a local recognizer produces results faster than a remote recognizer, but the results from the slower remote recognizer may be expected to be of higher quality than the local results.

[0026] FIG. 4 shows an embodiment for choosing a first or second result. The embodiment illustrated in FIG. 4 differs from the embodiment in FIG. 3 as follows: if the score returned from the first recognizer indicates high enough confidence, then there is no need to request a result from the second recognizer. A threshold value may be configured which, when crossed, indicates that the chooser may use the result from the first recognizer without getting a second opinion from another recognizer. In the description herein, we assume that better confidence in the returned result is indicated by a relatively higher score (that is, a score with a greater number). However, a person of ordinary skill in the art will appreciate that if a lower score indicates greater confidence, then accordingly, the threshold to cross will be configured to be a lower number.

[0027] Upon receiving a first score, the chooser compares the first received score to a high threshold in Step 41. If the score is above the high threshold, the embodiment chooses in Step 42 a first result associated with the first score as the basis for creating a response, without waiting for another result. If the first score is not above the high threshold, the embodiment sets a timeout duration as a function of the first score in Step 44. Upon receiving a second result before the timeout occurs in Step 45, the embodiment presumes that the second result is superior to the first result and chooses it in Step 47. If the timeout occurs in Step 45 before receiving a second result, the embodiment chooses the first result in Step 46 as the basis for creating a response.

[0028] FIG. 5 shows an embodiment for choosing a first or second result. The chooser receives a first score and sets a timeout duration as a function of the score in Step 54. If the timeout occurs in Step 55 before receiving a second result, the chooser chooses the first result in Step 56 as the basis for creating a response. Upon receiving a second score and result before the timeout occurs in Step 55, the chooser compares the second score to the first score in Step 58. If the second score is more than the first score, the chooser chooses the second re-

sult in Step 57 as the basis for creating a response. If the second score is not more than the first score in Step 58, the chooser chooses the first result in Step 56.

[0029] FIG. 6 shows an embodiment for choosing a first or second result. The embodiment illustrated in FIG. 6 determines how to proceed after receiving a score from the first recognizer by comparing the score to both a high threshold and a low threshold. The low threshold is used to determine if the score associated with the first result is so poor that the chooser will not use the first results under any circumstances. When the score is lower than the low threshold, the chooser has no confidence in the first result, and disregards the first result as a basis for creating the response.

[0030] Upon receiving a score associated with a first result, the chooser compares the score to a high threshold in Step 61. If the score is above the high threshold, the chooser chooses the first result as the basis for creating a response in Step 62 without waiting for a response from the second recognizer, or in some embodiments even without requesting such a response,.

[0031] If the score is not above the high threshold in Step 61, the chooser compares the score to a low threshold in Step 63. If the score is not below the low threshold in Step 63, the chooser sets a timeout duration as a function of the score in Step 64. When the score is between the low and high thresholds, the corresponding result may be considered as the basis for the response depending on the score received for results from the second recognizer, if results are received from the second recognizer before the timeout. Upon receiving a second result before the timeout occurs in Step 65, the chooser chooses assumes that the second result is more accurate than the first, and chooses it in Step 67. If the timeout occurs in Step 65 before receiving a second result, the chooser chooses the first result in Step 66 as the basis for creating a response.

[0032] If the score is below the low threshold, the embodiment ignores the associated result in Step 68. The chooser proceeds to set a pre-configured timeout in Step 69, and does not base the timeout duration on a function of the score. Upon receiving a second result before the timeout occurs in Step 610, the chooser chooses the second result in Step 67 regardless of the associated score. In another embodiment, if the score associated with the second result is below a low threshold for the second score, the chooser may produce no useful response and signals an error. If the timeout occurs in Step 610 before receiving a second result, the embodiment produces no useful response and signals an error in Step 611.

[0033] If the embodiment of FIG. 6 chooses the second result, it is because the second recognizer produced a result that the first recognizer could not. This is typically because the second recognizer has a larger vocabulary or more accurate language or acoustic model. Upon choosing the second result in Step 67, the embodiment proceeds to update the first recognizer in Step 612 with vocabulary, language, and acoustic models from the sec-

ond recognizer.

[0034] Some embodiments use multiple similar recognizers. Some embodiments use different kinds of recognizers. Some embodiments with different kinds of recognizers perform a step of normalizing scores from different recognizers before comparing the scores to thresholds or to scores from other recognizers. Scores are most often scalar. Various embodiments represent scores on linear, logarithmic, or other scales. Various embodiments base scores on hypothesis probability calculations of phonemes, phonetic sequences, n-grams, word sequences (such as transcriptions), grammatically correct sentences (such as parses), and recognized interpretations of utterances according to domains of knowledge. Some embodiments combine two or more ways of computing scores into a single scalar score. Some embodiments use multi-dimensional scores based on retaining two or more ways of computing scores.

[0035] FIG. 7A shows curves for six possible functions 701 used to set a timeout duration, given a first recognition score. All curves lie below a maximum allowable latency 702. In range 703, first recognition scores are above a high threshold, and so the embodiment immediately chooses the first result without waiting for another. In some embodiments, the high threshold is the maximum possible first recognition score, so that the chooser always waits for a second response for some time. In range 704, first recognition scores are below a low threshold, so the chooser ignores the first recognizer result and the timeout duration for receiving a response from the second recognizer is set to the maximum latency 702. If no second result arrives before the timeout occurs, the chooser produces no useful response and signals an error. In some embodiments, the low threshold is the minimum possible first recognition score (e.g. 0), so that the embodiment never ignores the first result as a possible basis for a response, even if the first result is extremely low scoring.

[0036] Various curves 701 take different shapes such as linear, parabolic, s-shaped, and staircase. All curves 701 are non-increasing, and most are decreasing, or step-wise decreasing.

[0037] FIG. 7B illustrates an embodiment that waits for a second response unless the score of the first result exceeds a threshold. The dashed line of curve 705 indicates the timeout value as a function of the first recognition score. For low values of the first recognition score, the timeout duration is set to a maximum latency value 702. When the first recognition score exceeded the high threshold 710, then the chooser immediately chooses the first result without waiting for another result, and the embodiment responds without waiting for a second result. This is the behavior illustrated in FIG. 4.

[0038] Some embodiments compute a recognition score, not directly from a hypothesis strength within a recognizer, but as a probability of a second recognition score being above a threshold of desired improvement over the first recognition score. In some embodiments,

the improvement threshold changes over time.

[0039] As will be apparent to practitioners of the art, descriptions herein can be extended to systems of more than two recognizers. Any plural number of recognizers can be considered in making decisions such as where to send speech, whether to choose a result immediately, discard the result, or wait for another, how to compare scores, and whether to start a timeout timer, and what function of one or more timeout timers to use are

Responding early to scores above a threshold

[0040] FIG. 8 shows a flow chart of choosing a result on which to base a response. Speech is sent to both first and second recognizers at substantially the same time in step 81. Next, the chooser receives a result and recognition score from one of the recognizers, and in step 82 compares the recognition score to a threshold. As mentioned earlier, different thresholds may be used depending on which recognizer produced the received recognition score, or the same threshold may be used for scores received from either of the recognizers. When the same threshold is used, the score may be normalized before comparing against the threshold. When the recognition score exceeds the appropriate threshold, the chooser chooses the result for response processing in step 84. This is useful in embodiments for which latency is important, and the fastest response possible is desirable, as long as it meets a minimum desirable quality. Many variations are possible of systems with such behavior.

[0041] FIG. 9 shows a flow chart of a specific embodiment of the flow depicted in FIG. 4 in which a request to the second recognizer is not sent until after receiving a result from the first recognizer. The system begins by sending speech to a first recognizer, but not a second recognizer in step 91. Next, the chooser receives a result and recognition score from the first recognizer, and in step 92 compares the recognition score to a threshold. If the recognition score exceeds the threshold, the embodiment chooses the result for response processing in step 94. If, in step 92, the recognition score does not exceed the threshold, only then does the embodiment send the speech to the second recognizer in step 93. This is useful in embodiments for which first recognizers operate fast enough to avoid significant latency, and for which transmission or processing costs for the second recognizer are high. Many variations are possible of systems with such behavior.

Timeout as a function of score

[0042] FIG. 10 shows a flow chart of a specific embodiment. It begins by sending speech to a first and second recognizer in step 101. Next, the embodiment receives a result and recognition score from one of the recognizers, and in step 105 sets a timeout duration as a function of the recognition score. This is useful in embodiments

that define a specific function of tradeoffs between the lower latency of a first response and the potential incremental quality improvements of a second response. Many variations are possible of systems with such behavior.

[0043] FIG. 11 shows a flow chart of a specific embodiment. It begins by sending speech to a first and second recognizer in step 111. Next, the embodiment receives a result and recognition score from one of the recognizers, and in step 115 sets a timeout duration as a function of the recognition score. If a timeout occurs before receiving a second result, the embodiment chooses the first result for response processing in step 116. Many variations are possible of systems with such behavior.

Real-time transcription

[0044] Some embodiments operate on continuous speech. Such embodiments, on an effectively continuous basis, re-compute or adjust recognition scores and start or disable new timeout timers. Some such embodiments have multiple timers that run simultaneously. In various embodiments, continuous operation effectively means repeating operations on a timescale that is imperceptible to users, such as less than a few hundred milliseconds.

[0045] Some such embodiments are systems that display a continuously updated transcription as a user speaks. It is desirable to update the transcription within a certain maximum latency, and as soon as possible if the accuracy is sufficient. If a recognition score from a faster, but less accurate, recognizer exceeds a threshold, then the system updates the transcription with that recognizer's result. If the score does not exceed the threshold then the system waits for a response from a more accurate, but slower, recognizer. Some such embodiments repeatedly send speech to both recognizers and start timers every 10 milliseconds, expecting new results with a latency of 30 to 500 milliseconds. Accordingly, the system will have multiple timers running simultaneously and can switch between the results of one recognizer and the other on any frame boundary.

[0046] Some embodiments tend to favor local recognition results for real-time transcriptions, but choose more accurate, remotely-processed results of delimited spoken utterances as the basis for responses. Some embodiments that process delimited spoken utterances only respond to complete commands; automobiles, vending machines, humanoid robots, and some personal assistants may depend on such embodiments.

Physical implementations

[0047] Dual mode speech recognition, as described herein, is embodied in methods, in machines, and in computer-readable media that store code that, if executed by one or more computer processors, would cause the computer processors to perform speech recognition accordingly.

[0048] Some embodiments are implemented in modular ways, and various such embodiments use combinations of hardware logic modules and software function modules. Various modular embodiments perform different necessary functions within different comparable modules. For example, some embodiments have a module for receiving speech from a user, sending speech to a first recognizer, a module for sending speech to a second recognizer, a module for receiving a recognition score, and a module for detecting a timeout, and a module for updating a speech recognition vocabulary.

[0049] FIG. 12 shows an embodiment that is a mobile phone 121 with a display 122 that shows a live transcription as text 123. The embodiment is shown in a scenario of live transcription of speech for the utterances, "what's the weather?", at about $\frac{3}{4}$ of the way through the utterance. The recognition scores is a transcription hypothesis score, which is above a threshold and, therefore, causes the transcription text to update.

[0050] FIG. 13 shows an embodiment that is a rack of server processors 131. Each server processor is connected to the Internet, and runs software that instructs the server to receive utterances from remote devices, perform recognition, and send recognition scores and results to the remote devices.

[0051] FIG. 14A shows a first view of a speech recognition earpiece 141. It comprises a power switch 142 to turn the device on and off; a battery charging port 143 to allow charging an internal batter; an ear hook 144 that can suspend the device comfortably over a user's ear; a microphone slot 145 for receiving speech; and a speaker 146 for providing spoken responses to the user. FIG. 14B shows a second view of the speech recognition earpiece 141 in which power switch 142, ear hook 144, and speaker 146 are visible. FIG. 14C shows a third view of the speech recognition earpiece 141 in which battery charging port 143, ear hook 144, and speaker 146 are visible. Also visible is mode button 147, which is allows a user to enable and disable speech recognition.

[0052] FIG. 15A shows a non-transitory computer readable medium, which is a magnetic computer disk 151. It stores code that, if executed by one or more processors, would cause the processors to perform dual mode speech recognition according to an embodiment. Magnetic computer disks are commonly used to store code for server processors. FIG. 15B shows a non-transitory computer readable medium, which is a Flash random access memory (RAM) chip 152. It stores code that, if executed by one or more processors, would cause the processors to perform dual mode speech recognition according to an embodiment. Flash memory is commonly used both in data centers and in mobile devices to store code for processors in system-on-chip devices. Practitioners will recognize that all computer-readable media, except for propagating signals, are possible embodiments.

[0053] FIG. 15C shows a system-on-chip package 153, with solder balls arranged for surface mounting the

package to a printed circuit board within a device. FIG. 15D shows the flat top side of the system-on-chip package 153.

[0054] FIG. 16 shows a block diagram of a system-on-chip 160. It comprises a cluster of computer processor (CPU) cores 161 and a cluster of graphics processor (GPU) cores. The processors are connected through a network-on-chip 163 to an off-chip dynamic random access memory (DRAM) interface 164 and Flash interface 165. System-on-chip 160 also has a display interface 166 and I/O interface module 167 coupled to the memory interfaces. The I/O interface enables touch screen interfaces, microphones, speakers, and USB devices, such as keyboards and mice, among others, to access the memory interfaces. System-on-chip 160 also comprises a network interface 168 to allow the processors to access the Internet through wired or wireless connections. By executing instructions stored in RAM devices through interface 164 or Flash devices through interface 165, the CPUs 161 and GPUs 162 perform dual mode speech recognition according to an embodiment.

[0055] FIG. 17 shows a block diagram of a server system 170. It comprises an array of CPUs 171 and an array of GPUs 172 connected through a board-level interconnect 173 to a RAM 174 and network interface 175. By executing instructions stored in RAM 174, the CPUs 171 and GPUs 172 perform dual mode speech recognition according to an embodiment.

[0056] The disclosure comprises the following items:

1. A computer-implemented speech recognition method comprising:

sending speech audio data to a first recognizer;
and
receiving, from the first recognizer, a first result and a first associated recognition score.

2. The method of item 1, comprising, if the first recognition score is above a threshold, responsive to a determination that the first recognition score is above the threshold, using the first result as a basis for creating a response to the speech audio data.

3. The method of item 1 or 2, comprising sending the speech audio data to a second recognizer.

4. The method of item 3, comprising receiving the speech audio data, wherein the steps of sending the speech audio data to the first and second recognizers both take place in response to receiving the speech audio data.

5. The method of item 3 or 4, comprising, if the first recognition score is below the threshold, responsive to a determination that the first recognition score is below the threshold, waiting for a period of time to receive, from the second recognizer, a second result

and a second associated recognition score.

6. The method of item 5, comprising, if the second result and the second associated recognition score are received from the second recognizer within the period of time, responsive to a determination that the second result and the second associated recognition score have been received from the second recognizer within the period of time, using whichever of the first or second result that has the higher respective associated recognition score as a basis for creating a response to the speech audio data. 5 10

7. The method of item 5 or 6, comprising, if the second result and the second associated recognition score are not received from the second recognizer within the period of time, responsive to a determination that the second result and the second associated recognition score have not been received from the second recognizer within the period of time, using the first result as a basis for creating a response to the speech audio data. 15 20

8. The method of any of items 5 to 7, wherein the length of the period of time is determined based on a function of the value of the first associated recognition score. 25

9. The method of item 8, wherein the function is such that the length of the period of time is inversely related to the first associated recognition score. 30

10. The method of any of items 3 to 9, wherein the step of sending the speech audio data to the second recognizer is responsive to a determination that the first associated recognition score is below the threshold. 35

11. The method of any of items 3 to 10, further comprising: 40

responsive to the first associated recognition score being below a lower threshold, ignoring the first result; and
responsive to not receiving a second response before a timeout occurs, signaling an error. 45

12. The method of any preceding item, wherein the first recognizer and the second recognizer are local. 50

13. The method of any of items 1 to 11, wherein the first recognizer and the second recognizer are remote.

14. The method of any of items 1 to 11, wherein the first recognizer is local and the second recognizer is remote. 55

15. The method of any preceding item, wherein the speech audio data is a continuous audio stream.

16. The method of any preceding item, wherein the speech audio data is a delimited spoken query.

17. The method of any preceding item, wherein the recognition score is based on a phonetic sequence score.

18. The method of any preceding item, wherein the recognition score is based on a transcription score.

19. The method of any preceding item, wherein the recognition score is based on a grammar parse score.

20. The method of any preceding item, wherein the recognition score is based on an interpretation score.

21. The method of any preceding item, the method comprising responding to the speech audio data based on the basis for creating the response to the speech audio data.

22. A non-transitory computer readable medium storing code that, when executed by one or more computer processors, causes the one or more computer processors to carry out the method of any preceding item.

23. A non-transitory computer readable medium storing code that, when executed by one or more computer processors, causes the one or more computer processors to:

send speech to a first recognizer and a second recognizer;
receive, from the first recognizer, a first result associated with a recognition score;
set a timeout duration based on a function of the value of the recognition score; and
responsive to receiving no result from the second recognizer before the timeout occurs, choose the first result as a basis for creating a response.

24. A method for performing dual mode speech recognition, comprising:

receiving a spoken query from a user;
processing the spoken query, including sending the spoken query to a local recognition system on a mobile device;
transmitting the spoken query to a remote recognition system via a communications link; and
setting a latency timer to a preset timeout value; in the event that the spoken query is not recog-

nized by either the local recognition system or the remote recognition system within the latency timer period, choosing the recognition failure as a final result;

in the event that the spoken query is recognized by both the local recognition system and the remote recognition system within the latency timer period, obtaining a recognition result and associated recognition score from both the local recognition system and the remote recognition system;

choosing the final result as the recognition result associated with the higher recognition score; in the event that the spoken query is recognized by only the local recognition within the latency timer period, obtaining a recognition result and associated score from the local recognition system; and

choosing the local recognition result as the final result;

in the event that the spoken query is recognized by only the remote recognition system within the latency timer period, obtaining a recognition result and associated score from the remote recognition system; and

choosing the remote recognition result as the final result;

taking action on behalf of the user based on the final result; and

in the event that the spoken query is recognized by the remote recognition system within the latency timer period, upon determining that the remote recognition result contains vocabulary information not contained within a client vocabulary maintained within the local recognition system, requesting that the remote recognition system update the client vocabulary.

25. The method of item 24, wherein the length of time of the preset timeout value is determined based on a function of the value of the associated score from the local recognition system.

26. The method of item 25, wherein the function is such that the length of the period of time is inversely related to the first associated recognition score.

27. A non-transitory computer readable medium storing code that, when executed by one or more computer processors, causes the one or more computer processors to carry out the method of any of items 24 to 26.

Claims

1. A computer-implemented speech recognition method comprising:

sending speech audio data to a first recognizer and a second recognizer;

receiving, from the first recognizer, a first result and a first associated recognition score;

setting a timeout duration based on a function of the value of the first associated recognition score; and

responsive to receiving no result from the second recognizer before the timeout occurs, choosing the first result as a basis for creating a response.

2. The method of claim 1, wherein the function is such that the timeout duration is inversely related to the first associated recognition score.

3. The method of claim 1, wherein the function is non-increasing.

4. The method of any preceding claim, wherein the first recognizer and the second recognizer are local.

5. The method of any of claims 1 to 3, wherein the first recognizer and the second recognizer are remote.

6. The method of any of claims 1 to 3, wherein the first recognizer is local and the second recognizer is remote.

7. The method of any preceding claim, wherein the speech audio data is a continuous audio stream.

8. The method of any preceding claim, wherein the speech audio data is a delimited spoken query.

9. The method of any preceding claim, wherein the recognition score is based on a phonetic sequence score.

10. The method of any preceding claim, wherein the recognition score is based on a transcription score.

11. The method of any preceding claim, wherein the recognition score is based on a grammar parse score.

12. The method of any preceding claim, wherein the recognition score is based on an interpretation score.

13. The method of any preceding claim, further comprising: responsive to receiving a second result from the second recognizer, updating a client vocabulary.

14. The method of any preceding claim, the method comprising receiving, from the second recognizer, a second result.

15. The method of any preceding claim, the method

comprising responding to the speech audio data based on the basis for creating the response to the speech audio data.

16. A non-transitory computer readable medium storing code that, when executed by one or more computer processors, causes the one or more computer processors to carry out the method of any preceding claim. 5
- 10
17. A mobile device enabled to perform dual mode speech recognition, the device comprising:
- a module for receiving speech from a user; 15
- a module for sending speech to a first recognizer; 20
- a module for sending speech to a second recognizer;
- a module for receiving a recognition score; and 25
- a module for detecting a timeout, the timeout duration being a function of the value of the recognition score,
- wherein the mobile device chooses a first result from the first recognizer if it does not receive a result from the second recognizer before the timeout occurs. 30
18. The mobile device of claim 17, wherein the first recognizer is local to the device and the second recognizer is remote from the mobile device. 35
19. The mobile device of claim 17 or 18, further comprising:
- a vocabulary update module enabled to update a local recognition vocabulary. 40
- 45
- 50
- 55

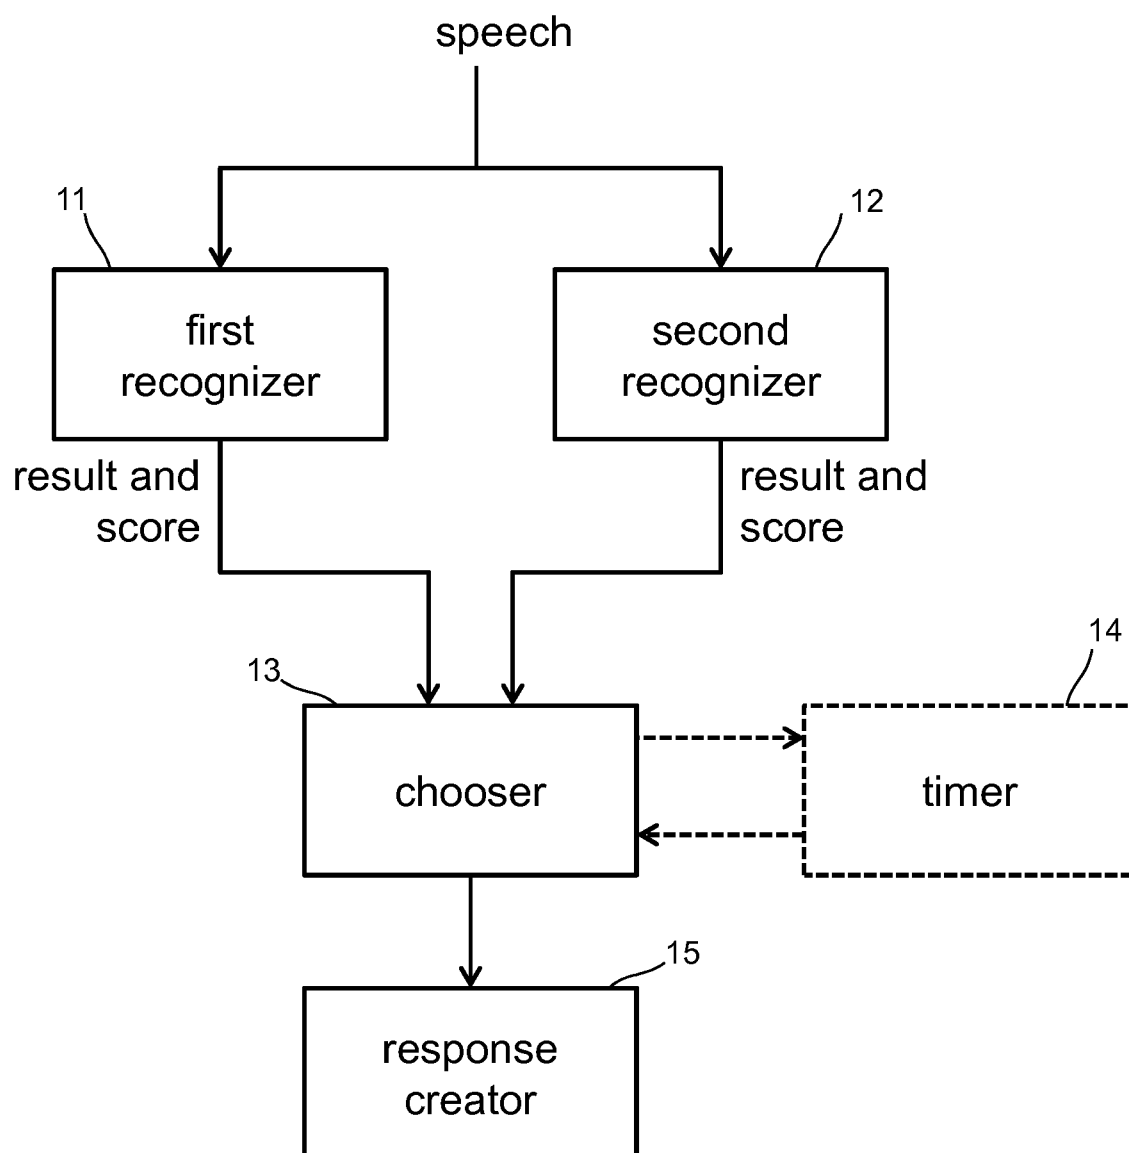


FIG. 1

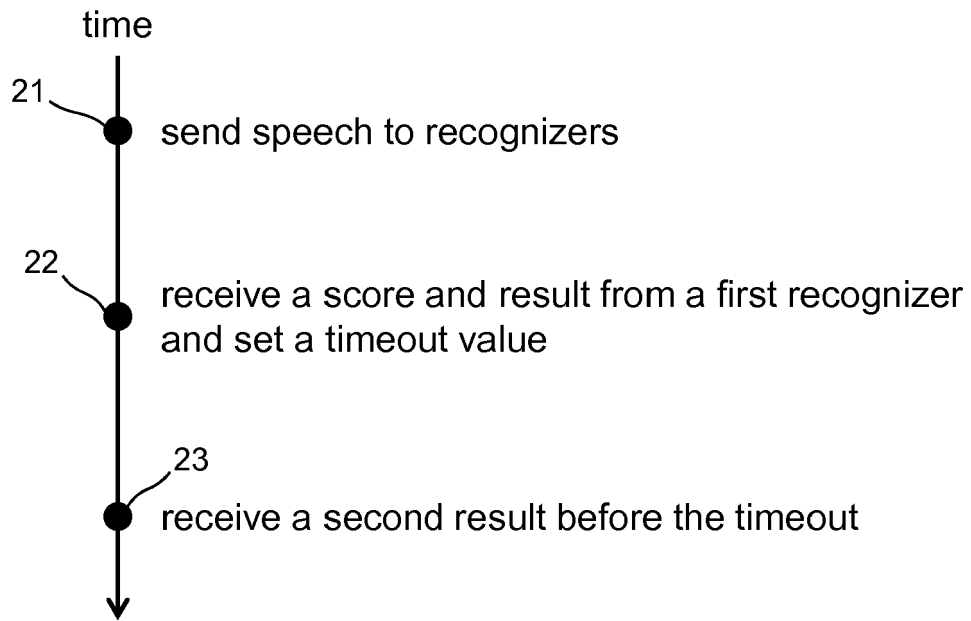


FIG. 2A

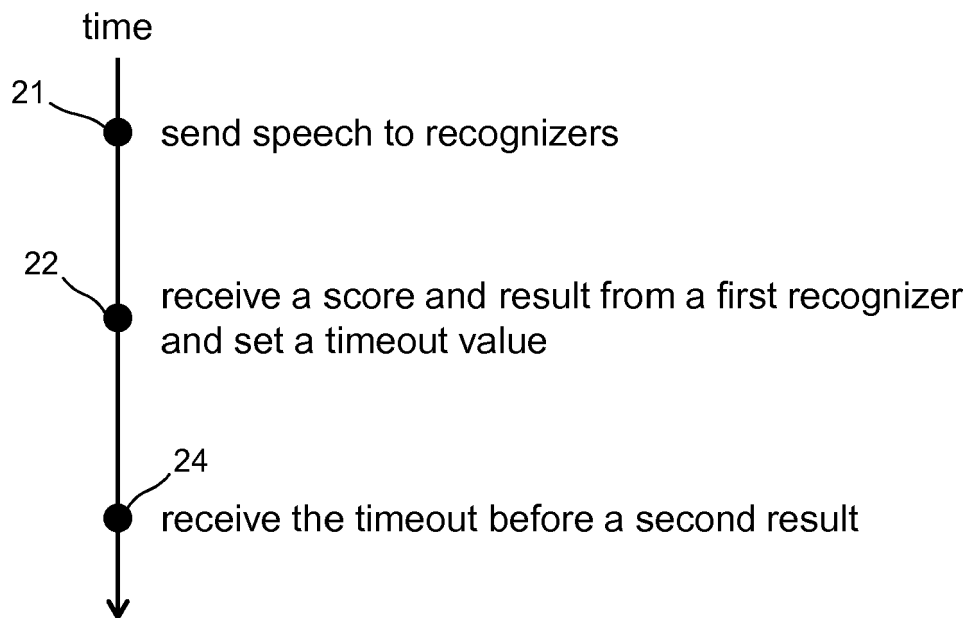


FIG. 2B




FIG. 3

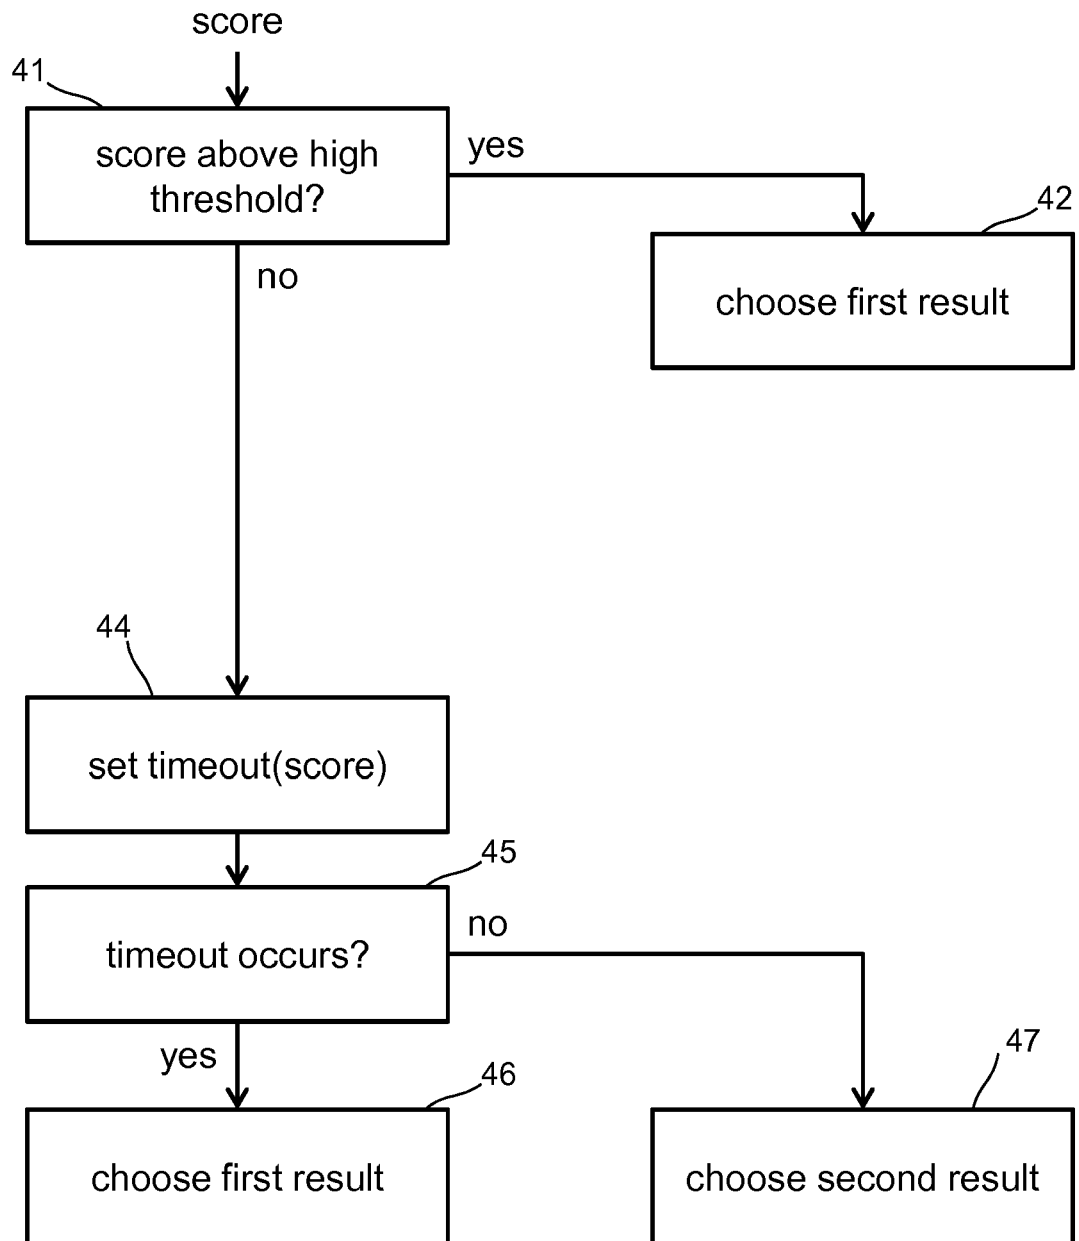


FIG. 4

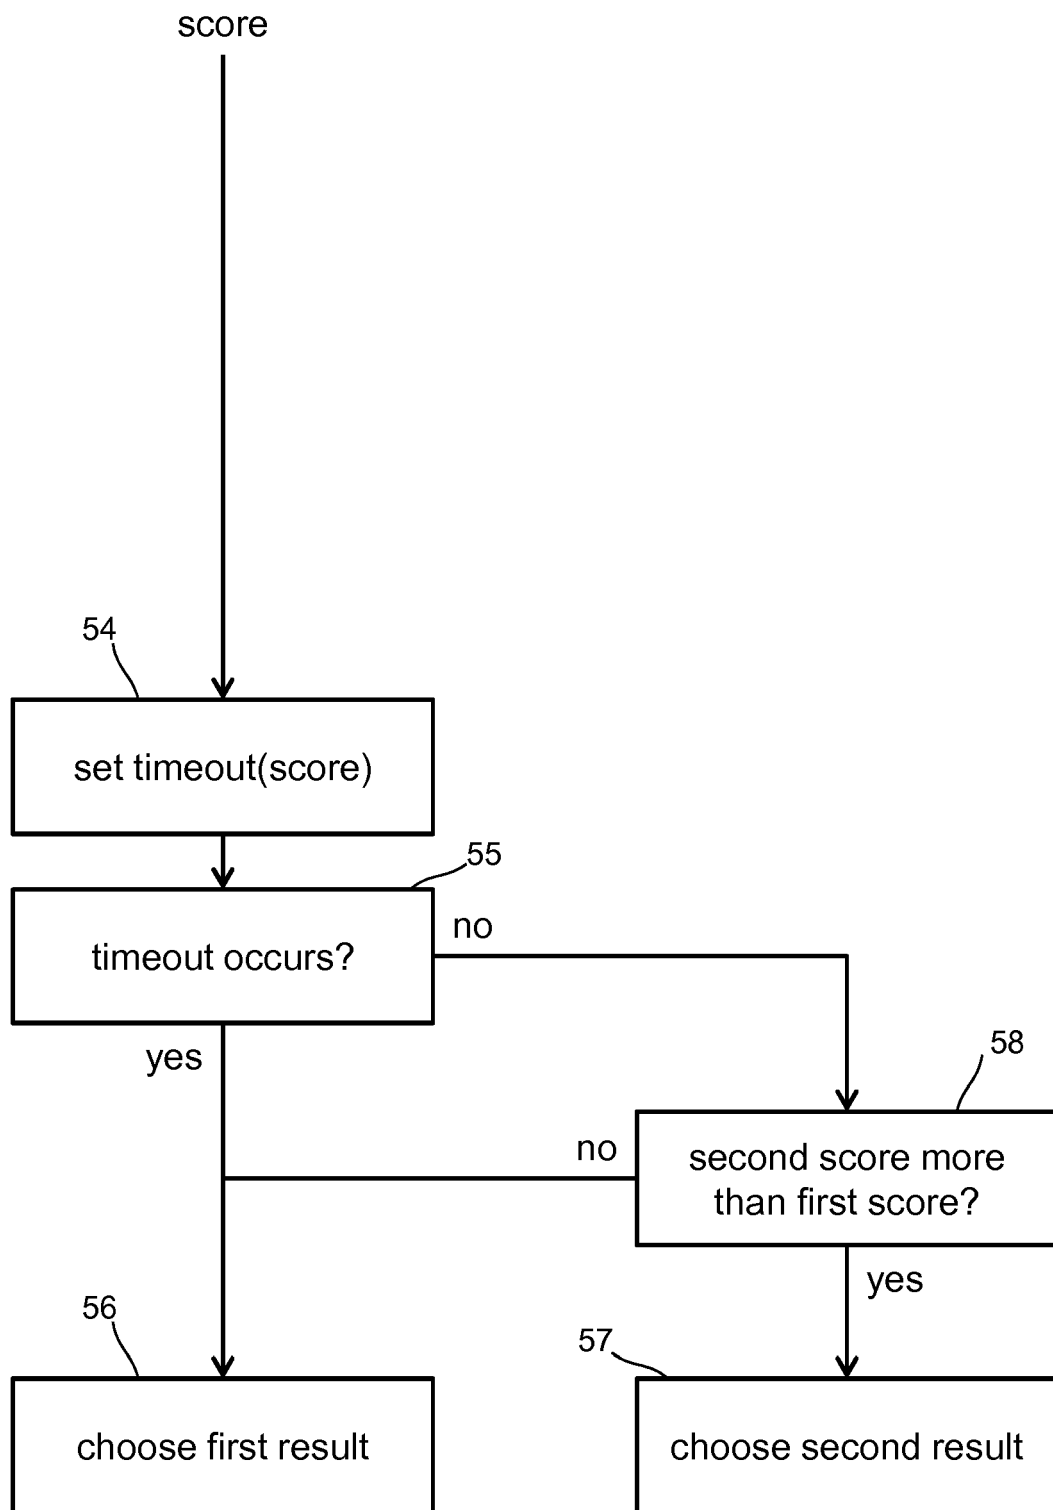


FIG. 5

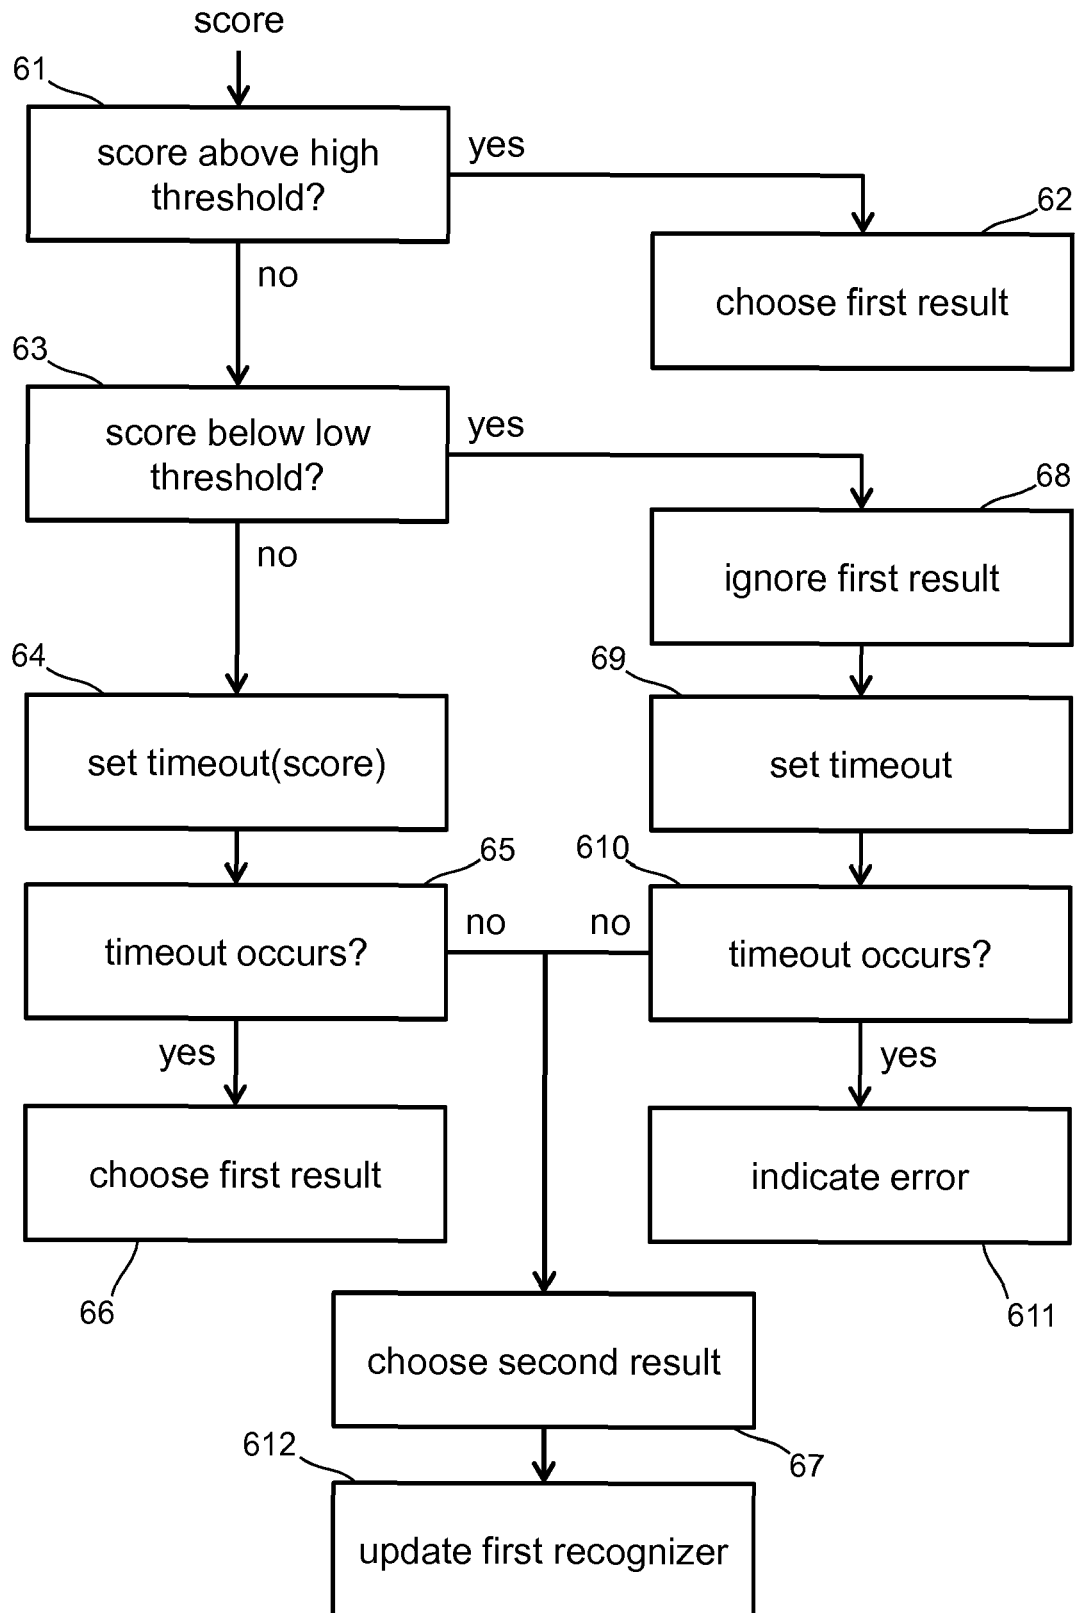


FIG. 6




FIG. 7A




FIG. 7B

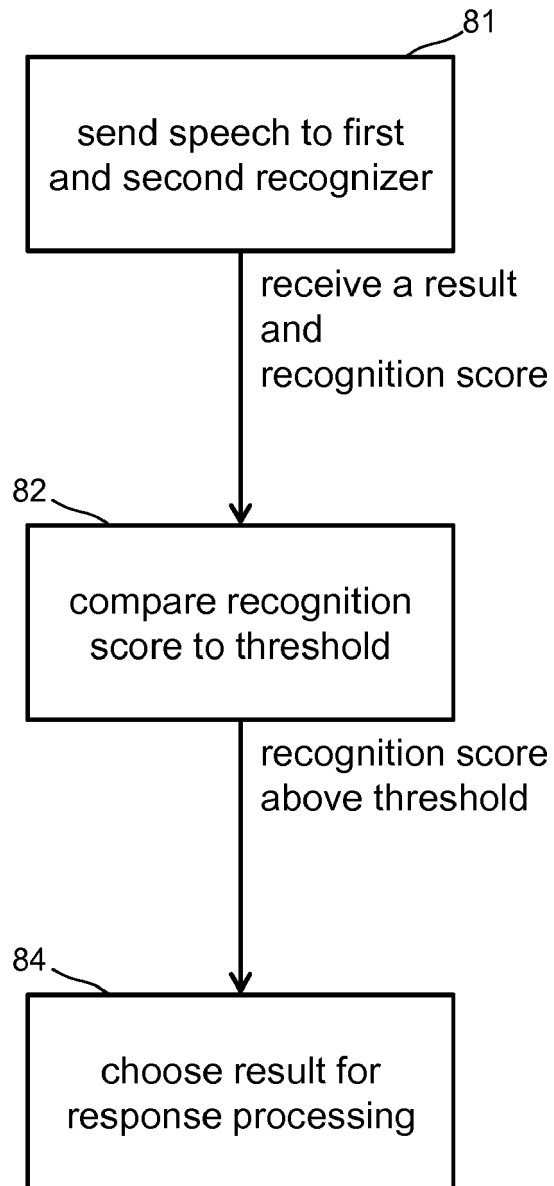


FIG. 8

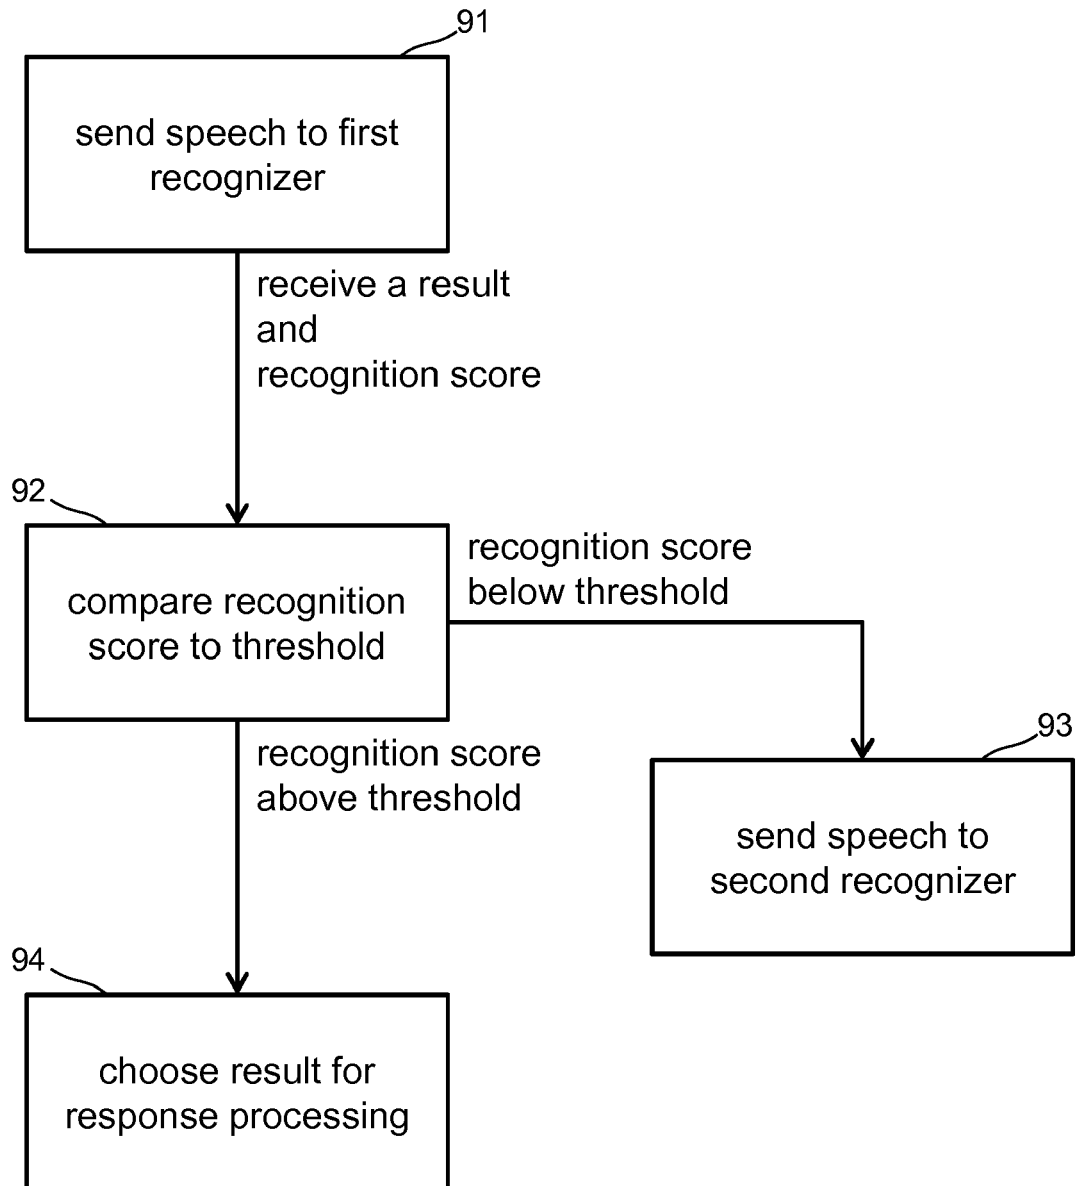


FIG. 9

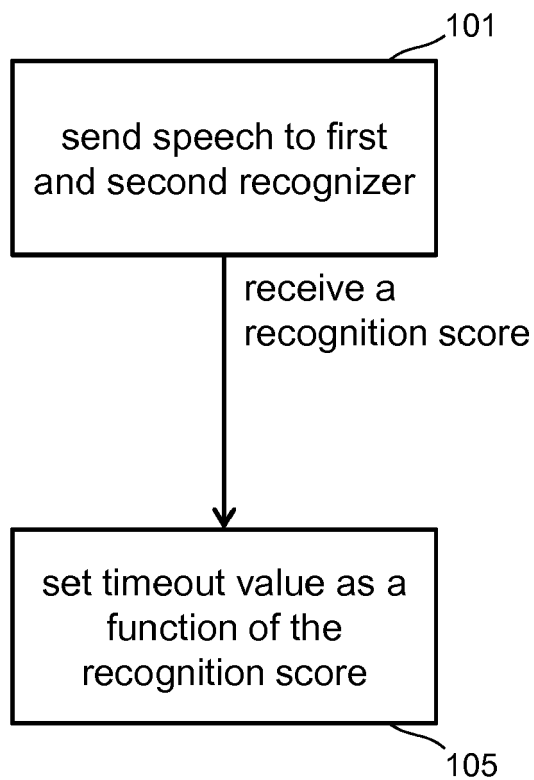


FIG. 10

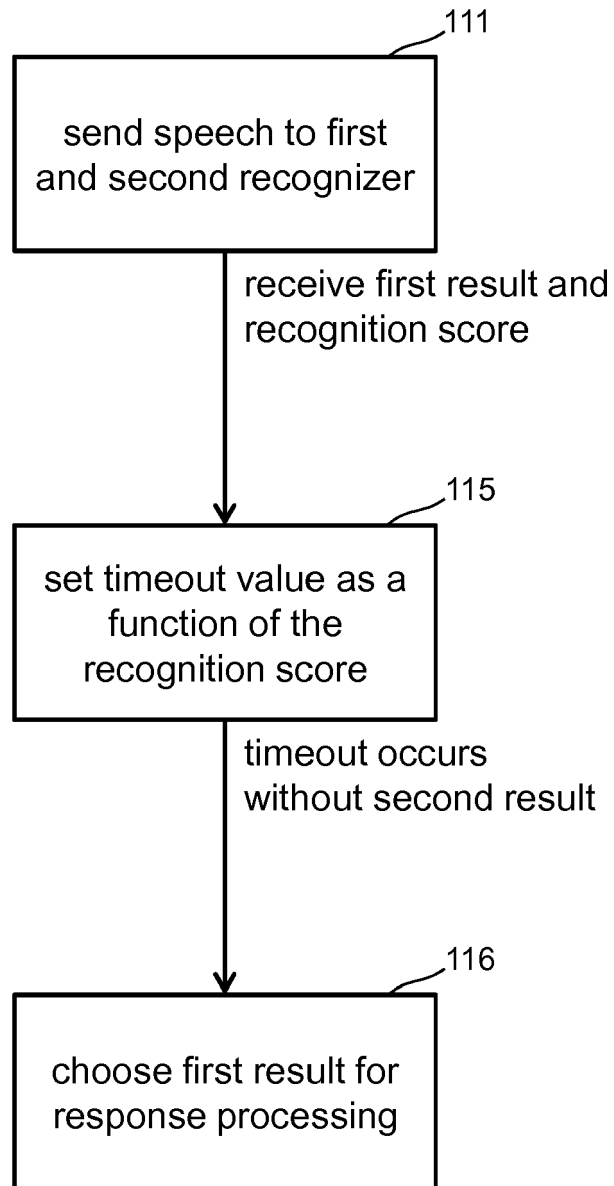


FIG. 11




FIG. 12




FIG. 13





FIG. 14C




FIG. 15A




FIG. 15B




FIG. 15C




FIG. 15D

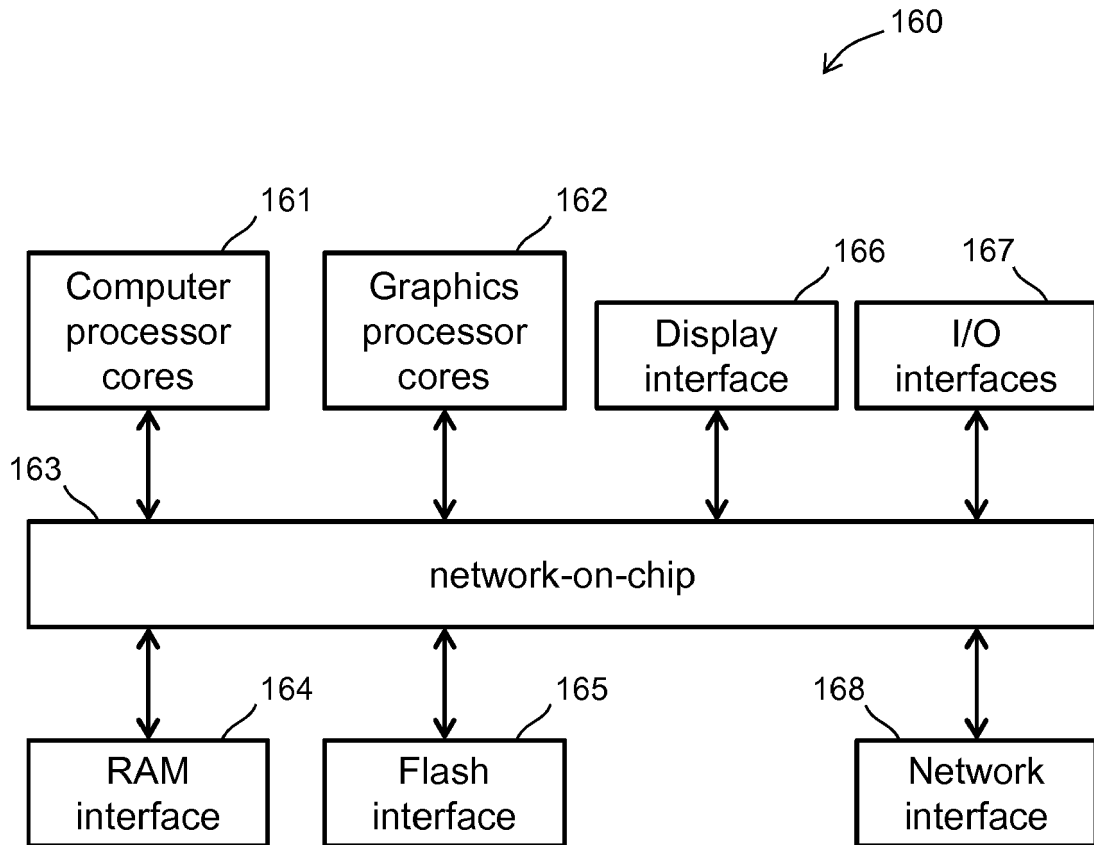


FIG. 16

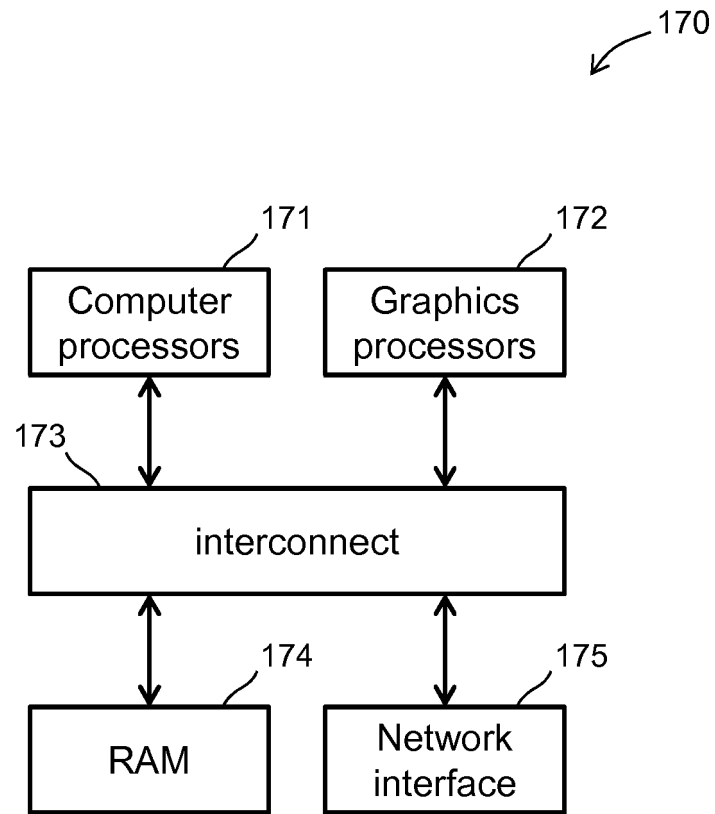


FIG. 17



EUROPEAN SEARCH REPORT

 Application Number
 EP 18 17 7044

5

10

15

20

25

30

35

40

45

50

55

9

EPO FORM 1503 03.82 (P04C01)

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (IPC)
X	US 2013/132084 A1 (STONEHOCKER TIMOTHY [US] ET AL) 23 May 2013 (2013-05-23) * abstract; figures 1-3 * * paragraph [0006] - paragraph [0007] * * paragraph [0028] - paragraph [0035] * * paragraph [0016] - paragraph [0018] * -----	1-19	INV. G10L15/30 G10L15/32
A	US 2004/210437 A1 (BAKER JAMES K [US]) 21 October 2004 (2004-10-21) * abstract; figures 1-3 * * paragraph [0007] - paragraph [0013] * * claims 1-17 * -----	1-19	
A	US 2012/179457 A1 (NEWMAN MICHAEL [US] ET AL) 12 July 2012 (2012-07-12) * abstract; figures 1-6 * * paragraph [0028] - paragraph [0061] * -----	1-19	
A	EP 2 930 716 A1 (SAMSUNG ELECTRONICS CO LTD [KR]) 14 October 2015 (2015-10-14) * abstract; figures 1-6 * * paragraph [0026] - paragraph [0074] * -----	1-19	TECHNICAL FIELDS SEARCHED (IPC)
A	WO 2016/209444 A1 (INTEL CORP [US]) 29 December 2016 (2016-12-29) * abstract; figures 1-6 * * page 3, line 25 - page 10, line 14 * -----	1-19	G10L
The present search report has been drawn up for all claims			
Place of search Munich		Date of completion of the search 15 August 2018	Examiner Képesi, Marián
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 18 17 7044

5

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

15-08-2018

10

15

20

25

30

35

40

45

50

55

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2013132084 A1	23-05-2013	US 2013132084 A1	23-05-2013
		US 2015154959 A1	04-06-2015
		US 2016217788 A1	28-07-2016
		US 2017256264 A1	07-09-2017

US 2004210437 A1	21-10-2004	NONE	

US 2012179457 A1	12-07-2012	US 2012179457 A1	12-07-2012
		US 2012179463 A1	12-07-2012
		US 2012179464 A1	12-07-2012
		US 2012179469 A1	12-07-2012
		US 2012179471 A1	12-07-2012

EP 2930716 A1	14-10-2015	CN 104978965 A	14-10-2015
		EP 2930716 A1	14-10-2015
		US 2015287413 A1	08-10-2015
		US 2017236519 A1	17-08-2017

WO 2016209444 A1	29-12-2016	CN 107660303 A	02-02-2018
		US 2016379626 A1	29-12-2016
		WO 2016209444 A1	29-12-2016
